

Abstract of the Disclosure

5

Methods, apparatuses and systems directed to a distributed data repository system including a plurality of symmetric data repository nodes. In certain embodiments of the present invention, the distributed data repository system is message-centric operative to store message payloads transmitted from client nodes. In certain embodiments, the distributed data repository system is BLOB-centric, maintaining binary data objects and indexes of attribute values that map to the binary data objects. Of course, the present invention can be utilized to store a great variety of digital data contained in message payloads. According to certain embodiments of the present invention, the attribute indexes are fully replicated across all data repository nodes, while the message payloads (e.g., data objects or other content) are
10 exchanged across data repository nodes as needed to fulfill client queries. In this manner, each data repository node in the distributed system can fulfill any client request, while reducing the storage and memory requirements for each data repository node. The reduced storage and computational requirements enable each distributed data repository node to be hosted by an inexpensive hardware platform and, therefore, allow for the deployment of large numbers of
20 distributed data repository nodes to achieve a distributed data repository system featuring high availability and reliability. In certain embodiments, each distributed data repository node is further equipped to act as an instant messaging (or other one-way messaging) server to allow client nodes to establish instant messaging connections with the data repository nodes in the distributed system.